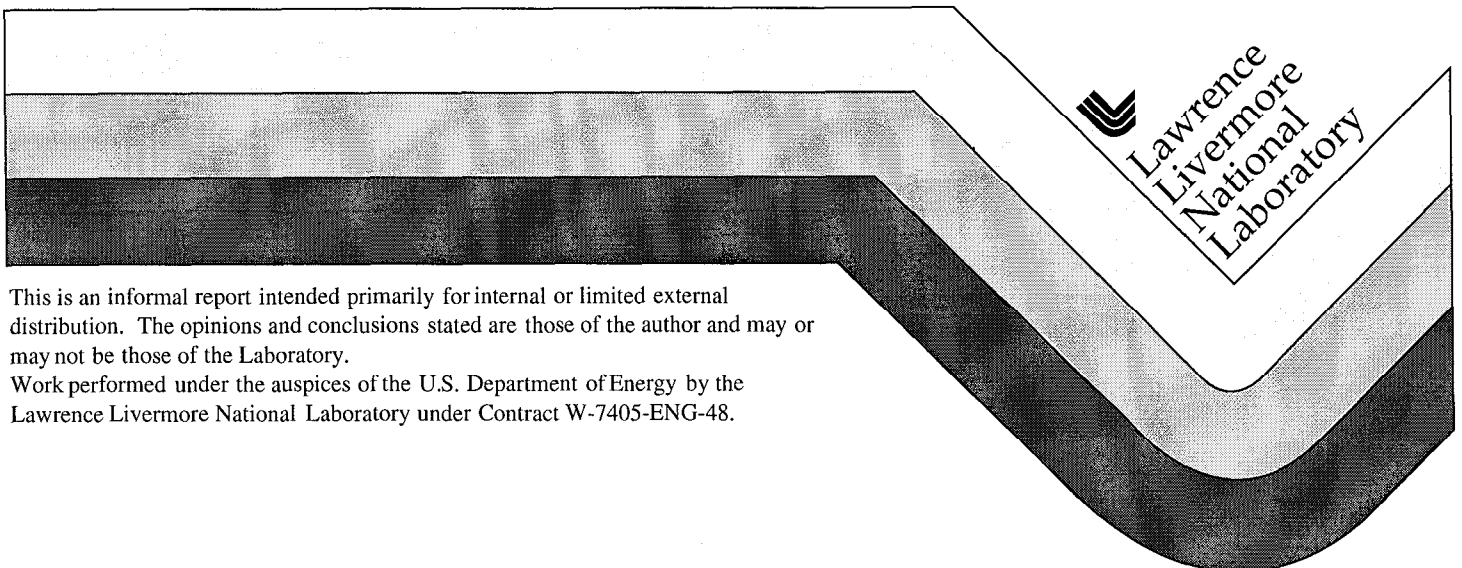


Image Recovery Techniques for X-ray Computed Tomography in Limited Data Environments

D. M. Goodman
J. A. Jackson
M. B. Aufderheide
E. M. Johansson

March 1, 1999



This is an informal report intended primarily for internal or limited external distribution. The opinions and conclusions stated are those of the author and may or may not be those of the Laboratory.

Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (615) 576-8401, FTS 626-8401

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Image Recovery Techniques for X-ray Computed Tomography in Limited Data Environments

Dennis M. Goodman (Laser Engineering Division)

Jessie A. Jackson (Laser Engineering Division)

Maurice B. Aufderheide (B Division)

Erik M. Johansson (Laser Engineering Division)

Abstract

There is an increasing requirement throughout LLNL for nondestructive evaluation using X-ray computed tomography (CT). In many cases, restrictions on data acquisition time, imaging geometry, and budgets make it unfeasible to acquire projection data over enough views to achieve desired spatial resolution using conventional CT methods. In particular, conventional CT methods are non-iterative algorithms that have the advantage of low computational effort, but they are not sufficiently adaptable to incorporate prior information or non-Gaussian statistics. Most currently existing iterative tomography algorithms are based on methods that are time consuming because they converge very slowly, if at all. The goal of our work was to develop a set of limited data CT reconstruction tools and then demonstrate their usefulness by applying them to a variety of problems of interest to LLNL. In this project we continued our development of reconstruction tools and we have demonstrated their effectiveness on several important problems.

I. Introduction:

Tomography is used throughout the laboratory to noninvasively investigate the composition and structure of objects. Conventional CT algorithms have the

advantage that are noniterative---hence they are computationally efficient---but they suffer from a lack of flexibility in that they can't incorporate prior information about the solution, nor can they use other than the squared error criterion to fit the data. Because the image recovery problem associated with most tomography applications is ill-posed, and this is particularly true for limited data situations, all possible prior information about the unknown object must be utilized in order to produce high-quality images. Furthermore, although least-squares (i.e. minimizing a sum of squares error function that measures the mismatch between the data and the model) is appropriate if the data has a Gaussian distribution with known constant variance, using it when the data has a non-Gaussian distribution can seriously degrade the quality of reconstructed images. An example is counting problems (e.g., emission tomography) where the data consists of particle or photon counts that typically have a Poisson distribution. Another example is outlier-corrupted data in which the distribution function is not known exactly, but it is known that the data pixels are subject to outliers that occur infrequently, but greatly distort those pixels where they occur. These outliers result from a variety of problems including a few bad detectors in a CCD array, improper assumptions about the model, or "hits" from extraneous radiation or particles. Outliers have highly deleterious effects on the reconstructed image when squared error is used.

The problems our techniques have been designed to handle can be described as follows. First there is a linear equation that models the relationship of the unknown to the data

$$\hat{\mathbf{y}} = \mathbf{A}\hat{\mathbf{x}} \tag{1}$$

where the vector $\hat{\mathbf{x}}$ represents the unknown image we wish to reconstruct, \mathbf{A} is a matrix that represents the effects of geometry, absorption, etc. that lead to the expected data $\hat{\mathbf{y}}$ if $\hat{\mathbf{x}}$ were the true image. The actual observed data is \mathbf{y} , and the

estimate of the true image is obtained by using an optimization algorithm to adjust $\hat{\mathbf{x}}$ so that $\hat{\mathbf{y}}$ is a best fit to the actual data according to some criterion function, say

$$f(\mathbf{y}, \hat{\mathbf{y}}) \quad (2)$$

where the choice of the criterion function, as we noted earlier, depends on the statistics of \mathbf{y} . Consequently, our reconstruction problem involves minimizing with respect to $\hat{\mathbf{x}}$ a function of the form

$$L(\mathbf{y}, \hat{\mathbf{x}}) = f(\mathbf{y}, A\hat{\mathbf{x}}) \quad (3)$$

This function usually is a negative log-likelihood function, so the image we recover is the maximum likelihood estimate of the unknown [1]. Because our problems typically are ill-posed, we add penalty parameters and impose prior constraints on \mathbf{x} , this leads to the minimization problem:

$$\hat{\mathbf{x}} = \underset{\hat{\mathbf{x}} \in S}{\text{Argmin}} \{ L(\mathbf{y}, \mathbf{x}) + \eta \|\mathbf{x}\|_1 + \lambda \|\mathbf{x}\|_2 \} \quad (4)$$

Here we are recovering a constrained penalized maximum likelihood estimate: the extra Euclidean norm and/or absolute value norm terms penalize large $\hat{\mathbf{x}}$, and the parameters η and λ determine the degree of penalty. It is also possible to include terms that penalize derivatives of \mathbf{x} , thereby imposing a larger penalty on its higher frequency components. If outliers are present, we do not use the maximum likelihood approach directly: instead, we select $f(\mathbf{y}, \hat{\mathbf{y}})$ to make the reconstructed image less sensitive to outliers. We give an example in the next section. The subset S represents upper and lower bounds on the components of the image vector $\hat{\mathbf{x}}$. These bounds provide crucial prior information about the solution and can greatly improve the quality of the recovered image. Because the unknown image usually represents nonnegative quantities such as absorption or energy, the most common constraint is that all of the components of $\hat{\mathbf{x}}$ be nonnegative. Other constraints on subsets of pixels can result from the known extent of the object and regions of

known voids or occlusions.

The optimization problem we have described above is a difficult one for several reasons. The most obvious is its very high dimensionality: both $\hat{\mathbf{x}}$ and \mathbf{y} frequently have 10^6 or more components. Another reason is the bounds that we have imposed on the solution require a constrained optimization technique. Many iterative non-linear minimization algorithms require storing an inverse of the matrix of second partials of $L(\mathbf{y}, \hat{\mathbf{x}})$ with respect to the components of $\hat{\mathbf{x}}$. This is clearly not possible for our problems because a 10^6 by 10^6 matrix is too large to store. The answer for unconstrained problems is either the conjugate gradient algorithm or the limited memory quasi-Newton algorithm; both of these algorithms in effect store low-order approximations of the second-order information that is provided by the inverse matrix of second partials.

For bound-constrained problems conventional optimization techniques [2-4] usually allow only one variable per iteration to attain a bound, so for very large problems these techniques are very slow because they spend too much time finding bounds. Consequently, standard iterative tomography algorithms [5-6] are based on methods that can attain multiple bounds in an iteration; unfortunately, however, they are essentially steepest descent techniques that use no second order information about $L(\mathbf{y}, \hat{\mathbf{x}})$ whatever, so they converge very slowly, if at all.

II. Progress:

II.a. Optimization Algorithm Development

Our ability to solve difficult tomography problems is the result of two specialized optimization algorithms that we have developed. The first is an extension of the conjugate gradient algorithm that incorporates bound constraints on the variables; it is described in [7]. This constrained conjugate gradient (CCG) algorithm is unique in

that it incorporates a bending linesearch that permits multiple bounds to be attained during a single iteration. It has previously been applied with great success to a variety of practical problems. These are described in Refs. [8-15]. The second is a limited-memory quasi-Newton algorithm that permits upper and lower bounds on the variables. We completed implementing our quasi-Newton (QN) algorithm as part of this LDRD project; it is based on the derivations in [16]. The very large dimensions of \mathbf{A} typically make calculating the matrix-vector product of Eq. 1 the most computationally expensive part of finding a solution, and which algorithm is most appropriate depends on the structure of \mathbf{A} .

In particular, the value of $L(\mathbf{y}, \hat{\mathbf{x}})$ and a related directional derivative must be calculated at each new sub-iteration within the bending linesearch. Completely recalculating Eq. 1 each time is very expensive, but in some cases a different approach is possible. This is because it is only necessary to recalculate that part of Eq. 1 corresponding to components of $\hat{\mathbf{x}}$ that have either attained or left bounds during a linesearch sub-iteration. For example, if only one component of $\hat{\mathbf{x}}$ has changed its bound status, then it is only necessary to compute one scalar-vector product---the product of that component with its corresponding column of \mathbf{A} ---rather than the entire vector-matrix product shown in Eq. 1. If \mathbf{A} has many columns, this saves considerable effort.

If the columns of \mathbf{A} are readily available, then CCG appears to be the best choice. However, in many tomography problems Eq. 1 is not computed in the usual manner. Often \mathbf{A} is a discretized version of the forward projection operator; therefore, it represents a set of line integrals, and it is very sparse. Consequently, Eq. 1 is computed via a *rule* that only calculates the line integrals and wastes no time on the large portions of \mathbf{A} consisting of zeros. Another example is where \mathbf{A} is the

kernel of a shift-invariant blurring function; in this case using a convolution algorithm based on the fast Fourier transform (FFT) is by far the most efficient method for computing Eq. 1. In some of our future work on the Advanced Hydrotest Facility (AHF) [17] we anticipate that \mathbf{A} will be of the form $\mathbf{A} = \mathbf{BC}$, where \mathbf{C} is a projection matrix and \mathbf{B} is a blur kernel that represents the effects of detector and source geometries. In this case both a sparse matrix rule and an FFT convolution algorithm are required. In all of the cases discussed in this paragraph, the programming effort required to calculate individual columns of \mathbf{A} is considerable, and often is not worthwhile. The advantage of the QN method is that bending is used first on a low-order quadratic approximation to $L(\mathbf{y}, \hat{\mathbf{x}})$ to produce a direction for the linesearch. This greatly reduces the need for bending during the linesearch itself, yet it still permits the algorithm to attain multiple bounds during an iteration. In the QN algorithm we have included the options both to perform additional bending in the linesearch if the columns of \mathbf{A} are available, and to perform no bending in the linesearch if they are not.

II.b Robust Tomography Algorithms.

As part of this project we completed work on and demonstrated the usefulness of our robust tomography technique. The squared error function is

$$f(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i=1}^N r_i^2 = \mathbf{r}'\mathbf{r} \quad (5)$$

Where the *residual* vector $\mathbf{r} = \mathbf{y} - \hat{\mathbf{y}}$ is one indication of mismatch between the observed data and the model of the data. If Eq. 5 is the criterion function in Eqs. 2-3, then the reconstructed image defined in Eq. 4 will be very sensitive to outliers. For example, suppose $N = 10,000$ in Eq. 5, and the magnitude of a typical component of

the residual is 1.0, except for one component r_k whose value is 100.0. Then r_k contributes roughly as much to the squared error as do all of the other components combined, and the minimization algorithm will try very hard to match the outlier y_k at the expense of all of the other data points. The solution is to use a *robust* criterion function [18-19] that reduces the influence of larger residuals. The most common robust criterion is

$$f(\mathbf{y}, \hat{\mathbf{y}}) = \sum_{i=1}^N g(r_i) \quad (6a)$$

where

$$g(r) = \begin{cases} r^2 & \text{if } |r| \leq c \\ 2c|r| - c^2 & \text{if } |r| > c \end{cases} \quad (6b)$$

The function $g(r)$ transitions smoothly from a squared penalty to a linear penalty, thereby reducing the influence of large residuals. The choice of c is data dependent; techniques for selecting it are given in [18-19].

We have demonstrated both that CCG is very effective in minimizing Eq. 4 when Eq. 6 is the criterion function and that this approach greatly improves image quality in the face of outliers. We present a simulated example. The true unknown is Fig. 1, and the resulting projection data with noise added is Fig 2. This data is in the form of a sinogram. A sinogram is a 2-dimensional image created by displaying all the ray sums at one angle versus all angles (projections) obtained. The noise is Gaussian mixture: with probability .99 the noise at a data pixel is Gaussian with standard deviation σ , but with probability .01 it is Gaussian with standard deviation 100σ . Consequently, there is an outlier 1% of the time. These outliers are obvious as speckles in Fig. 2. The reconstructed image using Eq. 5 is Fig 3; the reconstructed image using Eq. 6 is Fig 4. Outliers caused the streaks in Fig. 3. Our robust technique reduced the rms error between reconstruction and true object by a factor of two. Although robust statistical methods have been applied to a variety of problems, to

our knowledge this is their first application to tomography. The effectiveness of CCG made this possible. We have also successfully tried other robust criteria. Although we only presented a simulated example, many tomography problems are plagued by outliers, and this method has promise for a variety of real problems.

II.c Waste Drum Assay.

Using emission tomography to characterize mixed waste drums has been studied at LLNL for several years [20]. An active and passive computed tomography (A&PCT) technique has been developed that first uses an external radioactive source and active tomography to map the attenuation within a waste barrel. This attenuation map is used to define the matrix \mathbf{A} for the passive or emission tomography problem that is of interest to us here. Development of the A&PCT technique was done with program funding, but the emission tomography part of the problem gave us the opportunities to demonstrate the effectiveness of our techniques and to develop them for general emission tomography problems. At each detector position we acquire the entire gamma-ray spectrum and two counts of gamma-ray emissions are taken. The first count is in the region of a spectral peak of the isotope of interest; the second is in a region just outside this spectral peak. The purpose of the second measurement is to determine the level of background radiation and remove its effects on the first measurement. In previously developed algorithms, the net counts due to the isotope were obtained by subtracting the second measurement from the first. Two maximum likelihood expectation-maximization (MLEM) algorithms were developed, UCSF-MLEM [21] and APCT-MLEM [22], and applied to the corrected data to obtain a 3-D image of isotope activity. The sum of counts over all of the image voxels is related to an estimate of isotope activity within the drum. Unfortunately subtracting the two counts and then using a maximum likelihood

algorithm on the net counts is not a correct application of the likelihood principle. Furthermore, this approach can violate physical reality because there is a nonzero probability that a net count will be negative. The two MLEM algorithms lack the flexibility to implement the correct log-likelihood function. Using program funding, we derived the correct log-likelihood function [23]; as part of this LDRD project we developed a new algorithm for the waste drum problem, APCT-CCG, and we studied its behavior on both real and simulated data [24]. The flexibility of CCG made implementing the log-likelihood function relatively simple: it was only necessary to select the appropriate criterion function that accounted for Poisson statistics and incorporated peak and background measurements at each detector location. Fig. 5 compares applying the three algorithms to real data from known Pu-239 sources; APCT-CCG was out-performed on only one of the seven sources. Although the voxel sum is the most important parameter for waste drum assay, the reconstructed image is more important for most emission tomography problems, and APCT-CCG also did a much better job of reconstructing the image.

A test case was generated using simulated data. A simulated 3-D image was created with three slices. Each slice is 14 by 14 voxels. A point radioactive emission source of 30,350 counts was placed on the center slice at voxel location (5,5), i.e., just off the center of the slice. Using the system matrix the image was forward projected to create three sinograms. A level background equal to the maximum signal strength was added to these sinograms. The level was carefully chosen to be consistent with and representative of empirical data. The simulated sinograms were randomized by passing them through a Poisson random generator. Another set of sinograms were created with the same background level. These background sinograms were also randomized. Both the gross and background sinograms were used as input for the APCT-MLEM and APCT-CCG codes. The results are shown in Fig. 6. The APCT-

MLEM image is spread out over 3-by-3 voxels within each slice and across all three slices. Its total assay yields 36710 counts, i.e. 121% of the actual value. The APCT-CCG results are more localized and its assay value of 9936 counts is much closer to the original 30350 counts, i.e., within 1%. The other important observation is that the APCT-CCG code calculated sinograms are more representative of the original source data than the source plus background, which is not the case of the APCT-MLEM sinogram results. We also found this to be the case for real data as shown in ref [20].

For this problem it became obvious to us that model selection and validation are important issues, particularly when accurate estimates of parameters (e.g., isotope activity) are as critical as good images. In the case of Gaussian noise problems where the squared error function is used, tomography can be viewed as a linear regression problem. Consequently, the usual χ^2 tests can be applied to the squared error to measure goodness of fit and to choose between models, and the sinogram of the residuals is a good image to observe the adequacy of a particular model. However, this is not the case if the noise is not Gaussian. In studying the waste drum problem we realized that statistically it can be viewed as a *generalized linear model* in which there is a linear relationship between the unknown and the data, but the statistics are not necessarily Gaussian [19,25,26]. This is basically the model we defined in the Introduction. In this context it is possible to slightly modify $L(\mathbf{y}, \hat{\mathbf{x}})$ so that it becomes a *deviance function* that has the same minimum with respect to $\hat{\mathbf{x}}$, but exhibits behavior that is approximately χ^2 . Similarly, it is possible to define other sinogram images that behave as the usual residual image does for the squared error case. We have applied these ideas to the waste drum problem, and they will be useful for other tomography problems as well.

II.d. Neutron Imaging.

LLNL is currently developing a high-energy (10 - 15 MeV) neutron imaging system for use as an NDE tool in support of the Enhanced Surveillance Program (ESP). This approach to tomography promises to be a powerful technique for probing the internal structure of thicker objects that may be opaque to x rays and lower energy neutrons. Imaging experiments using neutron radiography were conducted at the Ohio University Accelerator Laboratory (OUAL) in FY98. The object being imaged was a right-circular PB cylinder with an outer diameter of 4" and a 2" diameter polyethylene insert, see Fig. 7. The insert was split into two half-cylinders with one serving as "blank" and the other having a series of 10-, 8-, 6-, 4-, 4--mm-diameter holes machined to depths of 0.5" into its outer (curved) surface. The areal density of the assembly ranged from 62.38 g/cm² (along the centerline) to 99.9 g/cm² (along the limb of the polyethylene insert). Reconstructions of this object using both filtered backprojection and our CCG algorithm with nonnegativity constraints on the reconstruction are shown in Fig. 6. The superiority of the CCG reconstruction is evident. Although we used the squared error criterion with CCG, in fact all of the raw data acquired during the OUAL experiments initially bore random sharp spikes rising several hundred to several thousand counts above the local average. These spikes were due to primarily to cosmic ray strikes in the CCD detector used to collect the data. Currently these spikes are removed by preprocessing prior to applying the tomography reconstruction algorithms. However, such data is a perfect candidate for the robust techniques we developed on this project, and we plan to apply them to neutron data in the near future. For a detailed description of the neutron imaging experiments, see Ref [27]

II.e The Advanced Hydrotest Facility.

In the first year of this project we adapted our CCG algorithm to cone beam tomography problems. In the second year we extensively used the resulting algorithm, CCG_Cone, for studies of limited view reconstruction at the Advanced Hydrotest Facility (AHF). This code has been an indispensable tool in these studies. It has been used to study the efficacy of reconstruction as a function of number of views, as well as how the orientation of views affects the quality of the reconstruction. These studies have allowed the AHF design group to make recommendations on how many views are needed, as well as where they should be placed. Currently, program funding is paying for the continued use of CCG_Cone on this problem. In addition, CCG_Cone has been used to study the effect of constraints on the quality of reconstructions. It has been found that a judicious use of constraints will help the AHF to achieve its objectives. More programmatic work will continue in this area.

We present an example using CCG_Cone. Our object consists of seven stacked slabs as shown in Fig. 9. Projection data, and reconstruction results are also shown in this figure. The origin of the axes is in the center of the middle slab. Four views are taken in the x-y plane. The angles of these views are equally spaced and are 0, 45, 90, and 135 degrees as measured in the x-y plane from the x axis. We used only in-plane viewing angles for purposes of comparison with the conventional non-iterative Feldkamp algorithm [28] that is only capable such angles. Experiments (not shown) with CCG_Cone on this example have demonstrated the advantage of using some out of plane viewing angles. The absorption of the slabs is arbitrarily set to 1.0 inside the slabs and 0.0 outside.

The large magnification and resulting large cone angle cause blurring at the extreme angles of the projections. Orthogonal slices through the reconstructions are displayed for both CCG_Cone and Feldkamp. Although a color scheme is used to display the reconstruction results for both algorithms, note that the scales for the two reconstructions are not the same. In particular, note that Feldkamp exhibits large excursions in both positive and negative directions. Since accurate spatial attenuation values are required in our image recovery, this clearly demonstrates the superiority of CCG_Cone.

IIf. Pulsed Photothermal Radiography.

In collaboration with researchers at several other institutions, we continued to apply our optimization algorithms to the problem of pulsed photothermal radiography (PPTR). This is essentially a tomographic method that inverts data from the time evolution of the heat equation, rather than from x-ray projection data, to see inside an opaque object. Our work on this project contributed to obtaining the first high-quality 3-D images of port-wine stain blood vessels [29-33]. Our CCG algorithm was used to solve the inverse problem to get these images. Furthermore, as part of this effort we demonstrated the feasibility of parallelizing our CCG algorithm [34]. An example of our work is shown in Fig. 10. This figure is a 1.6 mm x 1.6 mm x .6 mm volumetric image of *in vivo* port wine stain blood vessels of a human subject. The blood vessels are approximately 100 microns in diameter and are very well resolved. Given the highly ill-posed nature of heat equation inverse problems, this degree of resolution is highly impressive. This imaging technique is being used in a clinical trial at the UC Irvine Medical School to improve the outcome of laser treatment of port wine stains.

III. Future Work.

In future work we plan to continue refining our algorithms and applying them to practical problems. Immediate plans include parallelizing the algorithms, particularly the implementations of Eq. 1; further refinements of the forward projection model, including implementing the AHF blurring model; and further investigations of applying the generalized linear model formalism to tomography.

References

- [1] H. V. Poor, *An Introduction to Signal Detection and Estimation 2nd Ed.*, Springer-Verlag, 1994.
- [2] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization*, Academic Press, 1981.
- [3] R. Fletcher, *Practical Methods of Optimization, 2nd Ed.*, Wiley, 1987.
- [4] D. G. Luenberger, *Linear and Nonlinear Programming, 2nd Ed.*, Addison-Wesley, 1984. See pp 214-220.
- [5] H. Stark (Editor), *Image Recovery: Theory and Application*, Academic Press, 1987.
- [6] S. G. Azevedo, *Model-Based Computed Tomography for Nondestructive Evaluation*, PhD Thesis, LLNL Report UCRL-LR-106884, March, 1991.

- [7] D. M. Goodman, E. M. Johansson, and T. W. Lawrence, "On Applying the Conjugate Gradient Algorithm to Image Processing Problems," Chapter 11 in *Multivariate Analysis: Future Directions*, C. R. Rao, Editor, Elsevier Science Publishers, 1993.
- [8] D. M. Goodman, T. W. Lawrence, E. M. Johansson, and J. P. Fitch, "Bispectral Speckle Interferometry to Reconstruct Extended Objects from Turbulence-Degraded Telescope Images," Chapter 13 in *Handbook of Statistics, Vol. 10: Signal Processing and its Applications*, N. K. Bose and C. R. Rao, Editors, North Holland, 1993.
- [9] J. Kolman, W. S. Haddad, D. M. Goodman, and K. A. Nugent, "Application of a Constrained Optimization Algorithm to Limited View Tomography," *Proc. SPIE Conf.*, San Diego, July, 1994.
- [10] W. S. Haddad, J. E. Trebes, D. M. Goodman, H. R. Lee, I. McNulty, E. H. Anderson, A. O. Zalensky, "Ultra High Resolution Soft X-Ray Tomography", *Proc. SPIE Conf.*, San Diego, July, 1995.
- [11] J. R. Somoza, H. Szöke, D. M. Goodman, P. Béran, D. Truckses, S.-H. Kim, and A. Szöke, "Holographic Methods in X-ray Crystallography IV. A Fast Algorithm and its Application to Macromolecular Crystallography," *Acta Crystallographica*, Vol. A51, 1995.
- [12] T. E. Milner, D. M. Goodman, B. S. Tanenbaum, and J. S. Nelson, "Depth Profiling of Laser-Heated Chromophores in Biological Tissues by Pulsed Photothermal Radiometry," *Journal of the Optical Society of America-A*, Vol. 12 No. 7, July 1995.

[13] T. E. Milner, D. M. Goodman, B. S. Tanenbaum, B. Anvari, L. O. Svasand, and J. S. Nelson, "Imaging of Laser Heated Subsurface Chromophores in Biological Materials: Determination of Lateral Physical Dimensions," *Physics in Medicine and Biology*, Vol. 41, 1996.

[14] T. E. Milner, D. J. Smithies, D. M. Goodman, A. Lau, and J. S. Nelson, "Depth Determination of Chromophores in Human Skin by Pulsed Photothermal Radiometry," *Applied Optics*, June, 1996.

[15] M. J. C. van Gemert, J. S. Nelson, T. E. Milner, D. J. Smithies, W. Verkruijsse, J. F. de Boer, G. W. Lucassen, D. M. Goodman, B. S. Tanenbaum, L. T. Norvang, and L. O. Svaasand, "Non-Invasive Determination of Port Wine Stain Anatomy and Physiology for Optimal Laser Treatment Strategies," *Physics in Medicine and Biology*, Vol. 42, 1997.

[16] R. H. Byrd, J. Nocedal, and R. B. Schnabel, "Representations of Quasi-Newton Matrices and Their Use in Limited Memory Methods," *Math. Prog.*, Vol. 63, No. 2, Jan. 1994.

[17] "Advanced Hydrodynamic Radiography Technology Development Plan," Predecisional Draft, US Department of Energy, Defense Programs, Office of Research and Inertial Fusion, February 16, 1996.

[18] P. J. Huber, *Robust Statistics*, Wiley, New York, 1981.

[19] R. H. Myers, *Classical and Modern Regression with Applications*, PWS-Kent,

Boston, 1990. (See chapter 7).

[20] H. E. Martz, G. P. Roberson, D. C. Camp, D. J. Decman, J. A. Jackson, and G. K. Becker, *Active and Passive Computed Tomography Mixed Waste Area Final Report*, LLNL report, Nov., 1998.

[21] J. K. Brown, K. Kalki, J.A. Heanue, and B.H. Hasegawa, "Quantitative SPECT Reconstruction Using Multiray Projection Integrators," *Conference Record 1995 IEEE Nuclear Science Symposium and Medical Imaging Conference*, Vol. 2(1995)1272-1276.

[22] E. Keto, S. G. Azevedo, G. P. Roberson, D. J. Decman, H. E. Martz and E. M. Johansson, "Spatial Resolution Versus Signal to Noise in Quantitative Tomography" *Proceedings of the 4th Nondestructive Assay & Nondestructive Examination Waste Characterization Conference*, Salt Lake, Utah, October 24-26, 1995, pp. 405-420.

[23] D. M. Goodman, "Maximum Likelihood Estimation with Poisson (Counting) Statistics for Waste Drum Inspection," LLNL Report UCRL-ID-127361, May, 1997.

[24] J. A. Jackson, D. M. Goodman, G. P. Roberson, and H. E. Martz, "An Active and Passive Computed Tomography Algorithm with a Constrained Conjugate Gradient Solution," *Proceedings of the 6th Nondestructive Assay Waste Characterization Conference*, Salt Lake, Utah, November 17-19, 1998, pp. 325-358; UCRL-JC-130818, Lawrence Livermore National Laboratory, Livermore, CA, November 1998.

[25] P. McCullagh and J. A. Nelder, *Generalized Linear Models 2nd Ed.*, Chapman

and Hall, London, 1989.

[26] A. J. Dobson, *An Introduction to Generalized Linear Models*, Chapman and Hall, London, 1990.

[27] J. Hall, F. Dietrich, C. Logan, and G. Schmid, "Development of High-Energy Neutron Imaging in Support of Enhanced Surveillance Program Applications," LLNL Report, in preparation, 1999.

[28] L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical Cone-Beam Algorithm," *Journal of the Optical Society of America*, pp 612-619, June, 1984.

[29] T. E. Milner, S. A. Telenkov, B. S. Tanenbaum, J. S. Nelson, and D. M. Goodman., "Non-invasive Evaluation of Biological Materials Using Pulsed Photothermal Tomography," *Proc. The Winter Annual Meeting of ASME*, Anaheim, CA, November, 1998.

[30] S. A. Telenkov, D. M. Goodman, B. S. Tanenbaum, J. S. Nelson, and T. E. Milner, "Infrared Imaging of laser-heated portwine stains," Annual Meeting of the Optical Society of America, Baltimore, MD, October, 1998.

[31] S. Telenkov, D. J. Smithies, D. M. Goodman, B. S. Tanenbaum, J. S. Nelson, T. E. Milner, "Infrared imaging of in vivo microvasculature following pulsed laser irradiation," *Journal of Biomedical Optics*, pp 391-395, 1998.

[32] J. H. Torres, J. S. Nelson, B. S. Tanenbaum, T. E. Milner, D. M. Goodman, and B. Anvari, "Estimation of Internal Skin Temperatures in Response to Cryogen Spray

Cooling: Implications for Laser Therapy of Port Wine Stains," Submitted to *IEEE Journal of Selected Topics in Quantum Electronics*, Special Issue on Lasers in Medicine and Biology, 1999.

[33] S. A. Telenkov, B. S. Tanenbaum, D. M. Goodman, J. S. Nelson, and T. E. Milner, "In vivo Infrared Imaging of Laser-Heated Blood Vessels," Submitted to *IEEE Journal of Selected Topics in Quantum Electronics*, Special Issue on Lasers in Medicine and Biology, 1999.

[34] S. Shoari, N. Bagherzadeh, D. Goodman, T. E. Milner, D. J. Smithies, and J. S. Nelson, "A Parallel Algorithm for Pulsed Laser Infrared Tomography," *Pattern Recognition Letters*, (19) 5-6, pp 521-526, April, 1998.

Figure Captions:

Fig. 1 Original object.

Fig. 2 Sinogram of projections of object shown in Fig. 1 with Gaussian mixture noise added.

Fig. 3 Reconstruction from projection data in Fig. 2 using squared error criterion.

Fig. 4 Reconstruction from projection data in Fig. 2 using robust technique.

Fig. 5 Comparison of assay accuracy for three reconstruction methods.


Fig. 6 Simulated image and sinogram reconstruction.


Fig. 7 Neutron Imaging Object


Fig. 8 Reconstruction of object from neutron imaging data using filtered backprojection and CCG.


Fig. 9 Reconstruction of seven slabs from cone beam projection data. This figure demonstrates the superiority of CCG_Cone over the conventional Feldkamp algorithm.


Fig. 10 3-D Reconstruction of *in vivo* port wine stain blood vessels using pulsed photothermal tomography.

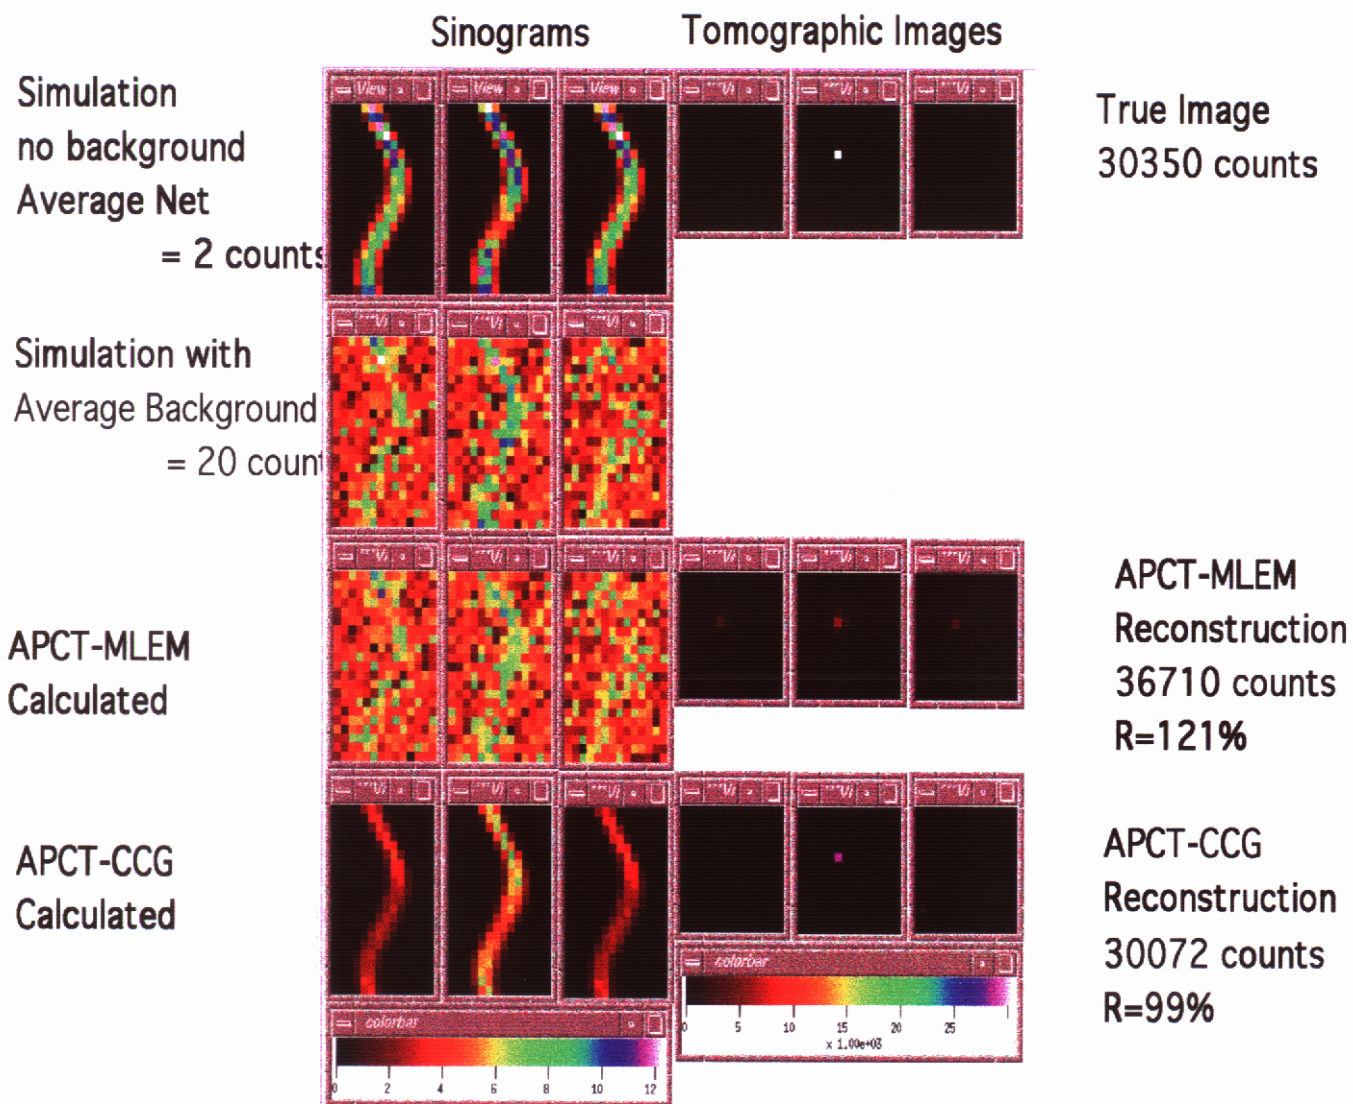


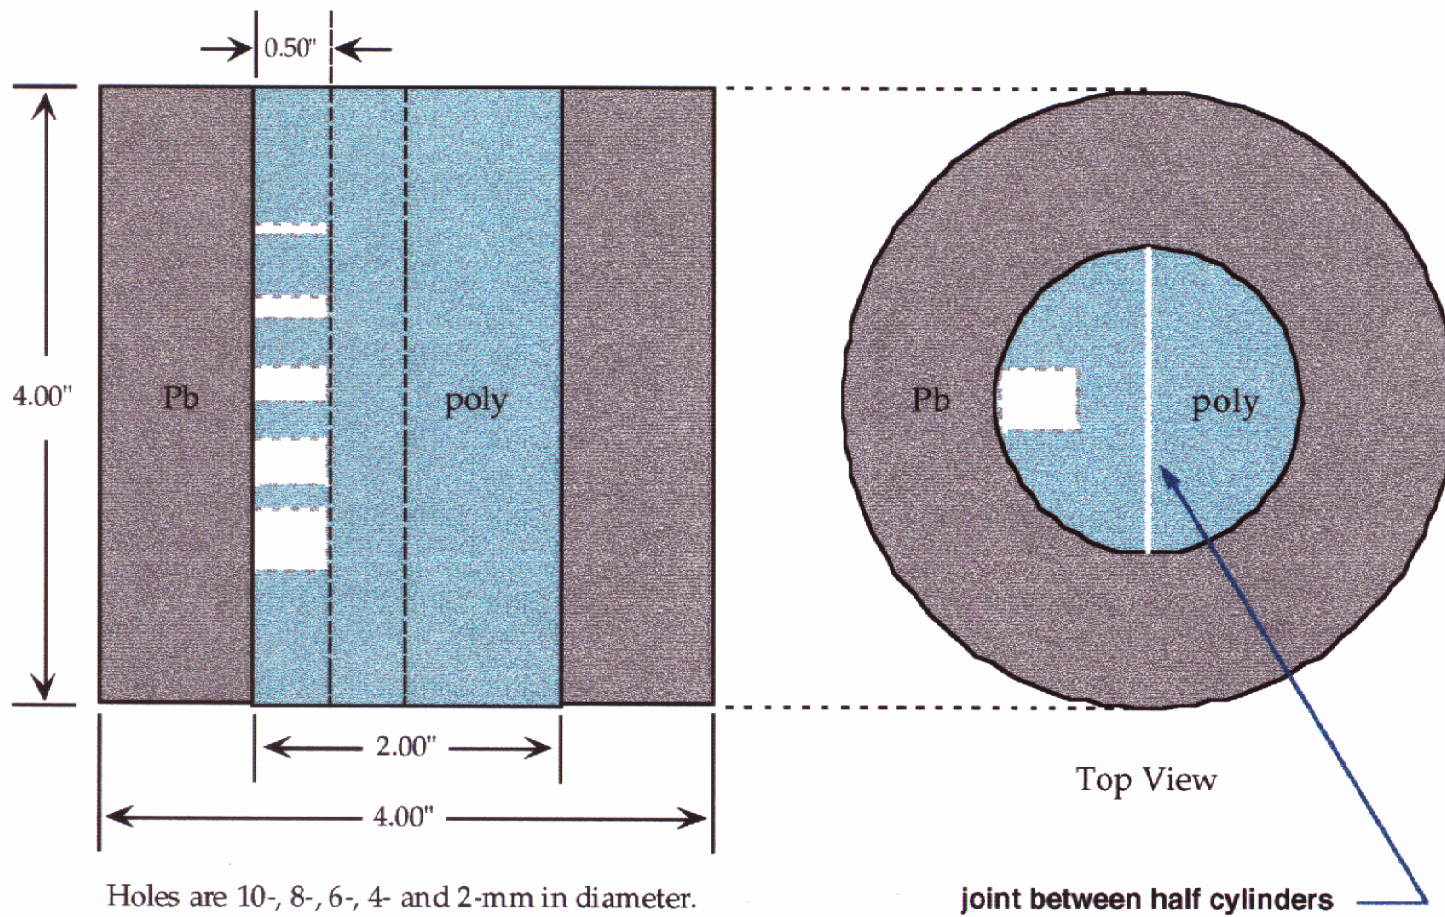

















Constrained Conjugate Gradient (CCG)



Standard Filtered Back Projection (FBP)



The CCG algorithm reduces pixel noise in reconstructed images by a factor of 3 - 5 compared to conventional FBP techniques

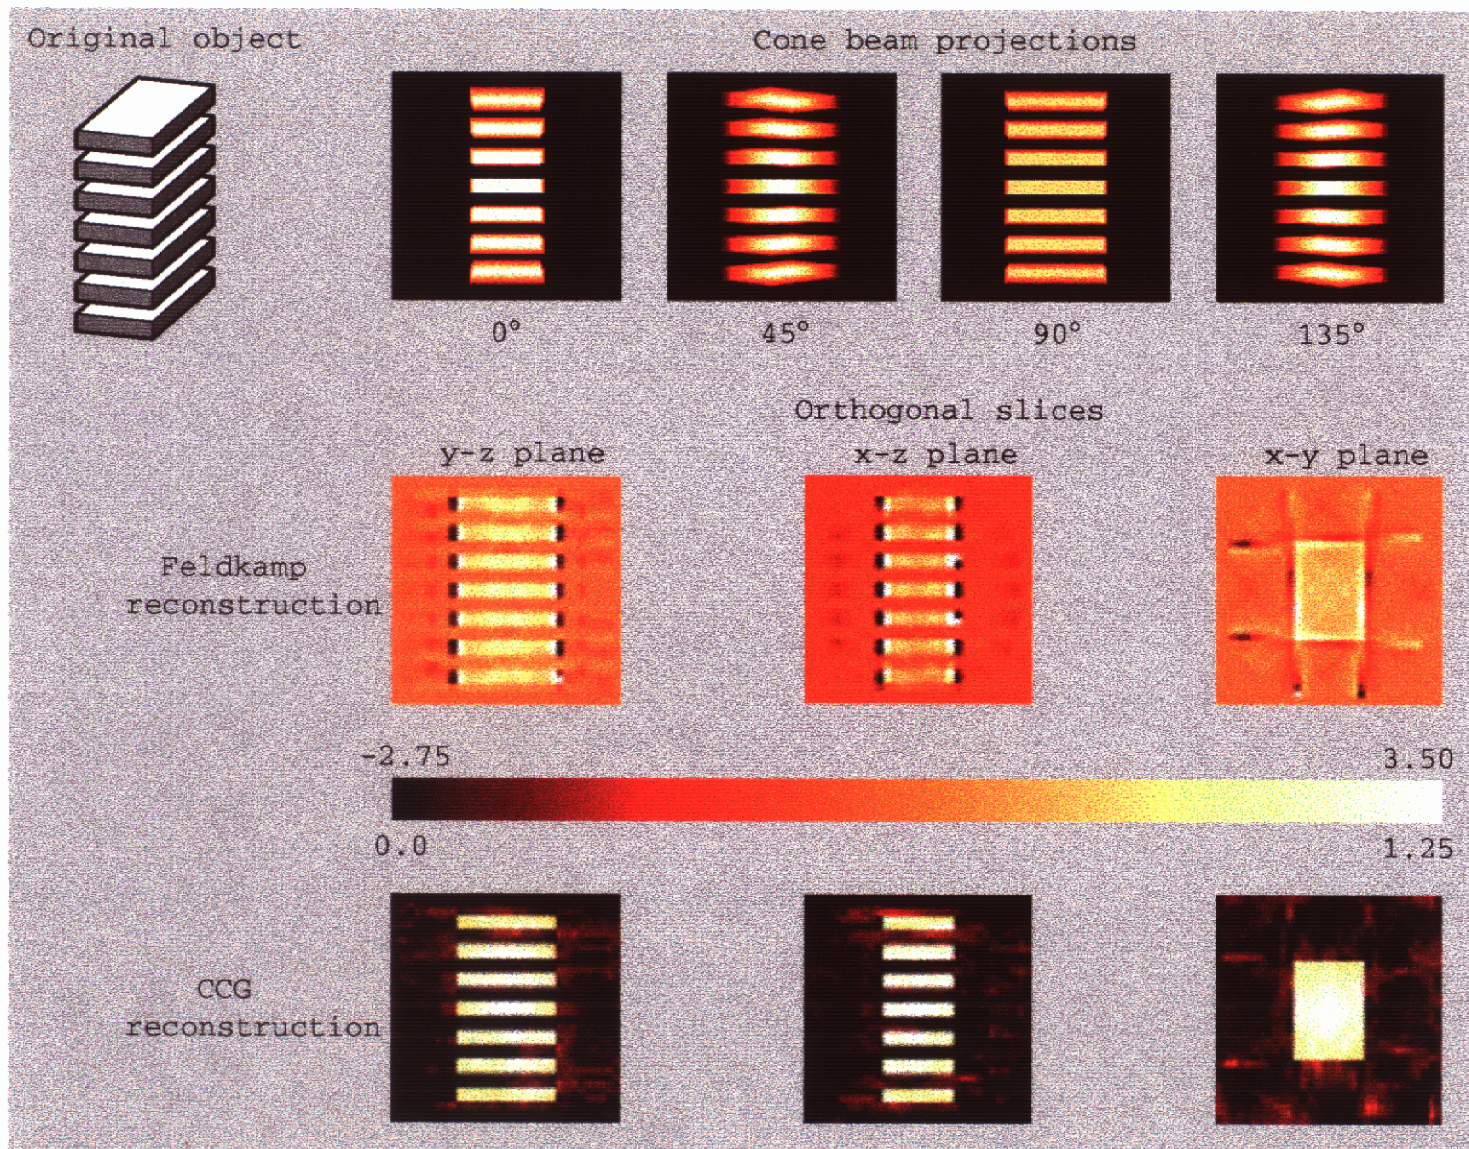


Fig01
TA705 Goodman/Martine

